of expertise. The output of the model is then compared with the performance of experts on a third task in which they judge the appropriateness of specific template features. Discrepancies are reduced by a "tuning process" until the match cannot be further improved.

The techniques employed in designing templates for diagnosis results in a representation of expertise at a fairly global level of knowledge and skill; they do not provide detailed information regarding the control processes through which experts are able to guide the course of diagnostic reasoning. These control processes incorporate highly refined heuristics about which the experts are almost wholly unaware. To discover the needed control knowledge we ask experts to complete tasks in which we have systematically perturbed aspects of the problem data. The data in these tasks are chosen so that members of an overlapping set of hypotheses will be elicited during the course of a problem solving episode. Eventual success depends on ability to overcome an initially plausible candidate in favor of a later, more correct alternative.

Results of the investigations to date reveal that problem solvers, even at the expert level, make frequent reasoning errors. Such errors are overcome either by the use of highly developed specific problem solving heuristics or by shifting to more detailed levels of reasoning. When our models are run on these same tasks, they initially make errors like the more expert subjects who are not able to shift levels in reasoning or otherwise recover from their initial (incorrect) hypothesis. By careful study of the behavior of experts who are able to solve such problems, a set of problem solving heuristics are identified. When these heuristics are incorporated into the model, it performs like the successful experts.

D. List of Relevant Publications

Connelly, D., & Johnson, P.E., Medical problem solving. Human Pathology, 1980, 11, (5), 412-419.

Elstein, A., Gorry, A., Johnson, P., & Kassirer, J. Proposed research efforts. In Clinical Decision Making and Laboratory Use. D. C. Connelly, E. Benson, & D. Burke (Eds.), University of Minnesota Press (in press).

Feltovich, P. J., Knowledge based components of expertise in medical diagnosis. Unpublished doctoral dissertation, University of Minnesota, 1981.

Feltovich, P. J., Johnson, P. E., Moller, J. H., & Swanson, D. B., The role and development of medical knowledge in diagnostic expertise. In W. Clancey & E. H. Shortliffe (Eds.), Readings in Medical AI, (in press)

Johnson, P. E. Cognitive models of medical problem solvers. In D. C. Connelly, E. Benson, & D. Burke (Eds.) Clinical Decision Making and Laboratory Use. University of Minnesota Press (in press).

Johnson, P. E., What kind of expert should a system be. The Journal of Medicine and Philosophy, (in press).

Johnson, P. E., Severance, D. G., & Feltovich, P. J. Design of decision support systems in medicine: Rationale and principles from the analysis of physician expertise. Proceedings of the Twelfth Hawaii International Conference on System Sciences, Western Periodicals Co., 1979, 3, 105-118.

Johnson, P. E., Duran, A., Hassebrock, F., Moller, J., Prietula, M., Feltovich, P., & Swanson, D. Expertise and error in diagnostic reasoning. Cognitive Science, 1981, 5, 235-283.

Johnson, P. E., Hassebrock, F., Validating computer simulation models of expert reasoning. In Proceedings of the Sixth European Meeting on Cybernetics and Systems Research, Vienna, Austria, April 1982, (in press).

Johnson, P. E., & Thompson, W. B. Strolling down the garden path: Detection and recovery from error in expert problem solving. Proceedings of the Seventh International Joint Conference on Artificial Intelligence, Vancouver, B. C., August 1981. pp. 214-217.

Moller, J. H., Bass, G. M., Jr., & Johnson, P. E. New techniques in the construction of patient management problems. Medical Education, 1981, 15, 150-153.

Swanson, D. B., Computer simulation of expert problem solving in medical diagnosis. Unpublished doctoral dissertation, University of Minnesota, 1978.

Swanson, D.B., Feltovich, P. J., & Johnson, P. E., Psychological analysis of physician expertise: Implications for the design of decision support systems. In D. B. Shires & H. Wold (Eds.), Medinfo77. Amsterdam: North-Holland Publishing Co., 1977, pp. 161-164.

E. Funding and Support

Work being done in scientific reasoning this past year has been sponsored under a current NSF (SE079-13036; 1981-82) grant to Paul Johnson. The work in Law is being supported by a grant from the NSF Law and Social Science Program (NSF SES-8111295; 1981-83). The work in medicine has been sponsored by the University of Minnesota Medical School and by NICHD (T36-HD-17151 and HD-01136; 1978-81) and NSF (NSF/BNS-77-22075; 1979-82) grants to the Minnesota Center for Research in Human Learning of which Paul Johnson is principal investigator.

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

A. Medical Collaborations and Program Dissemination via SUMEX

We have made use of the SUMEX system to consult with Dr. Benjamin Kuipers of Tufts University and Richard Keller of Rutgers University on issues relevant to the design of the Galen system. We have similarly consulted with Dr. N. S. Sridharan and Dr. L. Thorne McCarty of Rutgers University about our work in lawyer reasoning.

B. Sharing and Interactions with Other SUMEX-AIM Projects

Our work is complementary to many of the current projects supported on SUMEX-AIM. We will be investigating certain aspects of expert problem solving in order to develop better organizational and knowledge acquisition strategies. Such work requires that we be able to build upon the extensive experience in knowledge engineering within the SUMEX-AIM communities. Specifically, we first need to investigate the number of existing programs in order to determine the degree to which they satisfy the design goals which we will be establishing. We then hope to use the program construction tools that are available in order to build prototype systems to illustrate our ideas.

C. Critique of Resource Management

(None)

III. RESEARCH PLANS

A. Project Goals and Plans

Near term. Our research with the SUMEX-AIM resource has two components. First, we are involved in selective reimplementation of the diagnosis program, as described above. Second, we are making extensions to the original work of knowledge capturing, also described above. In the coming year, we plan to use the architecture of the Galen program to direct further studies of expertise in medical diagnosis, law and science. These studies are aimed at understanding the problem solving knowledge of individuals at three levels of expertise: novice, apprentice, and expert. This work should be of use in the development of CAI systems by providing hypotheses about the student model, as well as in the design of interfaces between expert systems and non-expert users.

In a related effort, we are studying the way in which experts determine the appropriate level of abstraction at which to analyze a particular problem. An important aspect of this process is to discover how the expert can recover from an initial incorrect choice of levels. We hope to determine the extent to which the difference between causal and prototypic modes of reasoning can be modeled as a difference in level of abstraction. If we are successful in this effort, we will be able to produce a common model for the two types of reasoning under study.

Long range. We propose to investigate the "knowledge capturing" process that occurs in the early stages of the development of expert systems when problem decomposition and solution strategies are being specified. Several related questions are being addressed: what are the performance consequences of different organization approaches, how can these consequences be evaluated, and what tools can assist in making the best choice? How can organizations be determined which not only perform well, but are structured so as to facilitate knowledge acquisition from human experts?

B. Justification and Requirements for Continued SUMEX Use

Our current model development is taking advantage of the sophisticated LISP programming environment on SUMEX. We have also benefited greatly from the interaction with other researchers facilitated by the SUMEX system. We expect to use SUMEX to allow other groups access to the Galen program. We hope in the future to make greater use of some of the knowledge engineering tools available on SUMEX.

C. Needs and Plans for Other Computing Resources Beyond SUMEX-AIM.

Once the Galen program has been fully implemented and tested, it will be ported to our VAX-11/780. The software is structured to allow simultaneous maintenance of both Interlisp and Franz Lisp versions. Computationally intense validation runs will be done on our VAX. SUMEX will continue to be used for collaborative activities and for program development requiring tools not available on the VAX.

D. Recommendations for Future Community and Resource Development

(None)

II.A.3 Pilot Stanford Projects

The following are descriptions of the informal pilot projects currently using the Stanford portion of the SUMEX-AIM resource pending funding, and full review and authorization.

II.A.3.1 Protein Secondary Structure Project

Protein Secondary Structure Project

Robert M. Abarbanel, M.D.
University of California Medical Center
University of California at San Francisco

I. SUMMARY OF RESEARCH

A. Project Rationale

Development of a protein structure knowledge base and tools for manipulation of that knowledge to aid in the investigation of new structures. System to include cooperating knowledge sources that work under the guidance of other system drivers to find solutions to protein secondary structure problems. Evaluations of structure predictions using known proteins and other user feedbacks available to aid user in developing new methods of prediction.

B. Medical Relevance and Collaboration

Many important proteins have been sequenced but have not, as yet, had their secondary or tertiary structures revealed. The systems developed here would aid medical scientists in the search for particular configurations, for example, around the active sites in enzymes. Predictions of secondary structure will aid in the determination of the full "natural" configuration of important biological materials. Development of systems such as these will contribute to our knowledge of medical scientific data representation and retrieval.

C. Highlights of Research Progress

The prediction of beta-alpha protein structures is nearly complete. A system has been developed on a VAX 11/750 at the University of California, San Francisco, to allow researchers to describe patterns of amino acid residues that will be sought in the sequences under study. The presence or absence of these "primary" patterns are then combined with other measures of structure, like hydrophobicity, to suggest possible alpha helix or beta sheet or turn configurations.

The segments of a sequence between turns are then analyzed to determine the allowable extent of the possible secondary structure assignments. Any segments remaining are then used to generate all possible complete structures. Only two beta strands with the character of sheet edges are allowed in any prediction. This heirarchical generation and pruning results in nearly 95% turn prediction accuracy, and excellent delimiting of helices and sheets. In some cases, one and only one secondary structure are predicted.

Research in Progress -- Funding of this research has been difficult during the year. An all volunteer effort to date has produced some very impressive results. Should some funding be acquired soon, work on other classes of proteins will continue. The original system written in the C language, will be constructed with a good scientist-user interface using facilities at SUMEX-AIM. The scheme of building structures by refinement of regions and higher level views of larger and larger sub-structures, will be expanded with a knowledge-engineered blackboard model. Preliminary design of these tools is under way.

D. List of Relevant Publications

The first paper on beta-alpha-beta protein structures is in preparation.

E. Funding Support

New Investigator Award approved but not funded by the National Library of Medicine, N.I.H.

II. INTERACTIONS WITH SUMEX-AIM RESOURCE

A. Medical Collaborations

None.

B. Sharing and Interactions with SUMEX Projects

This project is closely allied with the MOLGEN group, both in computer and scientific interests. Some pattern matching methodology created for the protein data base has been adopted and used in the various DNA knowledge bases. The principal persons in the MOLGEN group have contributed to this project's use and understanding of knowledge base software and resources.

C. Critique of Resource Management

The system load at SUMEX-AIM was so overwhelming so as to discourage the use of the DEC-10 system for development of tools. All programming and structure investigation has been done on the UNIX (TM) system at the University of California, San Francisco.

Resource management remains excellent. The staff are friendly and responsive. They cannot, however, despite large personal efforts, make up for the sorry lack of computing power that has plagued this project. It is hoped that improvements in load or hardware will allow continued development of the SUMEX-AIM software during 1982-83.

III. RESEARCH PLANS

A. Project Goals and Plans

Design of rules modules in UNITS editor that will allow an inexperienced user to express algorithms for structure prediction in near-natural language. These prediction schemata will then be translated into appropriate combinations of invocations of knowledge sources, editors, and feedback tools to aid the user in further refinement of his algorithms.

Systems to be developed for the discovery of rules and/or algorithms that can transform a protein sequence into a secondary structure, possibly with implications about tertiary structure. This will involve an effort along the lines of the meta-DENDRAL project.

Expansion of techniques used for beta-alpha prediction to other classes of proteins. Improvement of user interfaces to allow use of this sequence analysis system for problems of homology and energetics.

B. Need for Resources

1. SUMEX Resources

The availability of UNIX (TM) under SUMEX-AIM control will greatly aid in the transferability of existing algorithms. The environment of knowledge base tools and people is the primary motive for doing this work using SUMEX. Access to both established and developing systems aids this project in setting down standards of excellence, forward thinking about computing tools and methodologies, and active exchange of techniques and ideas. The close collaboration with the MOLGEN researchers is particularly useful in this regard.

2. Other Computing Resources

A soon to be established network connection with the Computer Graphics Laboratory at UCSF will provide access to 1) the latest in protein structural information, and 2) color line drawing graphics facilities for evaluation and display of this projects product. A real time display using color graphics will become a possibility.

Recommendations -- First and most important -- EXPAND the computing power available to SUMEX users. Facilitate networking with other computing environments like the Computer Graphics Laboratory at UCSF so that protein structural information may be exchanged and their hardware for 3d structure display may be utilized as a part of a complete biological structures analysis system.

II.A.3.2 Ultrasonic Imaging Project

ULTRASONIC IMAGING PROJECT

James F. Brinkley, M.D.
W.D. McCallum, M.D.
Depts. Computer Science, Obstetrics and Gynecology
Stanford University

I. SUMMARY OF RESEARCH PROGRAM

A. Project Rationale

The long range goal of this project is the development of an ultrasonic imaging and display system for three-dimensional modelling of body organs. The models will be used for non-invasive study of anatomic structure and shape as well as for calculation of accurate organ volumes for use in clinical diagnosis. Initially, the system has been used to determine fetal volume as an indicator of fetal weight; later it will be adapted to measure left ventricular volume, or liver and kidney volume.

The general method we are using is the reconstruction of an organ from a series of ultrasonic cross-sections taken in an arbitrary fashion. A real-time ultrasonic scanner is coupled to a three-dimensional acoustic position locating system so that the three-dimensional orientation of the scan plane is known at all times. During the patient exam a dedicated microcomputer based data acquisition system is used to record a series of scans over the organ being modelled. The scans are recorded on a video tape recorder before being transferred to a video disk. 3D position information is stored on a floppy disk file. In the proposed system the microprocessor will then be connected to SUMEX where it will become a slave to an AI program running on SUMEX. The SUMEX program will use a model appropriate for the organ which will form the basis of an initial hypothesis about the shape of the organ. This hypothesis will be refined at first by asking the user relevant clinical questions such as (for the fetus) the gestational age, the lie of the fetus in the abdomen and complicating medical factors. This kind of information is the same as that used by the clinician before he even places the scan head on the patient. The model will then be used to request those scans from the video disk which have the best chance of giving useful information. Heuristics based on the protocols used by clinicians during an exam will be incorporated since clinicians tend to collect scans in a manner which gives the most information about the organ. For each requested scan a prototype outline derived from the model will be sent to the microcomputer. The requested scan will be retrieved from the video disk, digitized into a frame buffer, and the prototype used to direct a border recognition process that will determine the organ outline on the scan. The resulting outline will be sent to SUMEX where it will be used to update the model. The scan requesting process will be continued until it is judged that enough information has been collected. The final model will then be used to determine volume and other quantitative parameters, and will be displayed in three dimensions.

We believe that this hypothesize verify method is similar to that used by clinicians when they perform an ultrasound exam. An initial model, based on clinical evidence and past experience, is present in the clinician's mind even before he begins the exam. During the exam this model is updated by collecting scans in a very specific manner which is known to provide the maximum amount of information. By building an ultrasound imaging system which closely resembles the way a physician thinks we hope to not only provide a useful diagnostic tool but also to explore very fundamental questions about the way people see.

We are developing this system in phases, starting with an earlier version developed at the University of Washington. During the first phase the previous system has been adapted and extended to run in the SUMEX environment. Clinical studies have been initiated to determine its effectiveness in predicting fetal weight and left ventricular volume. At the same time computer vision techniques are being studied in order to develop the system further in the direction of increased applicability and ease of use. We thus hope to develop a limited system in order to demonstrate the feasibility of the technique, and then to gradually extend it with more complex computer processing techniques, to the point where it becomes a useful clinical tool.

B. Medical Relevance

This project is being developed in collaboration with the Ultrasound Division of the Department of Obstetrics at Stanford, of which W.D. McCallum is the head.

Fetal weight is known to be a strong indicator of fetal well-being: small babies generally do more poorly than larger ones. In addition, the rate of growth is an important indicator: fetuses which are "small-for-dates" tend to have higher morbidity and mortality. It is thought that these small-for-dates fetuses may be suffering from placental insufficiency, so that if the diagnosis could be made soon enough early delivery might prevent some of the complications. In addition such growth curves would aid in understanding the normal physiology of the fetus. Several attempts have been made to use ultrasound for predicting fetal weight since ultrasound is painless, noninvasive, and apparently risk-free. These techniques generally use one or two measurements such as abdominal circumference or biparietal diameter in a multiple regression against weight. We recently studied several of these methods [McCallum 79] and concluded that the most accurate were about +/-200 gms/kg, which is not accurate enough for adequate growth curves (the fetus grows about 200 gms/week). The method we have developed is based on the assumption (which we have shown to be correct this past year [Brinkley 82d]) that fetal weight is directly related to volume since the density of fetal tissue is nearly constant. We have shown this year [Brinkley 82e] that by utilizing three dimensional information more accurate volumes and hence weights can be obtained.

In addition to fetal weight, the current implementation of this system has been evaluated for its ability to determine other organ volumes in vitro. In collaboration with Dr. Richard Popp of the Stanford Division of Cardiology we have evaluated the system on in vitro kidneys and latex molds of the human left ventricle [Brinkley 82b]. Left ventricular volumes are routinely obtained by means of cardiac catheterization in order to help characterize left ventricular function. Attempts to determine ventricular volume using one or two dimensional information from ultrasound has not demonstrated the accuracy of angiography. Therefore, three-dimensional information should provide a more accurate means of non-invasively assessing the state of the left ventricle.

C. Highlights of Research Progress

During the past year we have completed most of the evaluation of the first phase of our implementation and have analyzed and submitted the results for publication. The results have also been presented at three separate conferences. As mentioned in previous reports the first phase involves the manual outlining of scans with a light pen and the interpolation of an ad hoc model to the data. The system as it stands now is entirely data driven and hence has no notion of what kind of organ it is looking at. However, our results suggest that the data driven approach will still give more accurate volumes than any now obtainable, and thus will provide the medical impetus for continued research into knowledge driven approaches. The following evaluations have been completed:

1. <u>Engineering evaluations</u>: [Brinkley 82a], [Brinkley 82b] describe the implementation of the first phase of the system as well as the engineering evaluations. The evaluations included the reproducibility of three dimensional point determination and the accuracy of volumes computed on progressively more irregular objects. The reproducibility of 3D point determination was found to be about 6 mm. The various sources of this error were analyzed and found to come mostly from the ultrasound resolution of about 3 mm, and the position locator resolution of about 3 mm.

   Thirty volume trials on 10 water filled balloons gave 27 out of 30 calculations within 1.8 percent of true volume. Eighteen trials on 6 kidneys gave 17 out of 18 calculations within 5.1 percent of true volume. Fifteen trials on 5 human left ventricular molds gave 13 out of 15 calculations within -5.9 percent of true volume.

2. <u>Fetal weight estimation in vitro</u>: [Brinkley 82c],[Brinkley 82d] describe an evaluation of the ability of ultrasonic three-dimensional head and trunk reconstructions to estimate fetal weight on 25 dead neonates imaged in a water bath. The correlation between head plus trunk volume and measured weight was r=.985 with a standard error of 190 grams, over a weight range of 364 to 3650 grams. In this study we also established the empirical relationship between weight and volume: the correlation was r=.999, with a standard error of 37 grams.

3. Fetal weight estimation in utero: [Brinkley 82e] is in the final draft stages. It describes the results from our clinical study to determine the ability of volumes found by three dimensional ultrasonic head and trunk reconstructions to estimate in utero fetal weight. 41 live term fetuses were imaged within 48 hours prior to delivery. Head and trunk volumes calculated by the reconstruction technique were compared with measured birthweight in multiple linear regressions. In addition to the reconstruction volumes the three dimensional coordinates of various landmarks on the fetus were obtained with the system. The 3D distances between these landmarks were used to create additional predictor variables such as lengths, diameters, circumferences and simple volumes. The correlation between head and trunk volumes by reconstruction and the natural log of birthweight was r=.892, SE=91 g/kg. The best set of predictor variables was found by stepwise regression to be a head volume model as the product of three distances in the head, a trunk volume model as the product of three distances in the trunk, and the trunk volume by reconstruction. The correlation between these 3 predictors and the natural log of birthweight was r=.941, SE=69 g/kg. For comparison the accuracy reported by most methods in the literature have been about 100 g/kg. We have thus improved the clinical prediction of fetal weight by about 30 percent over current methods.

4. Conclusions from the evaluations of the data driven system: The in vitro results have shown that, in spite of a fairly poor 3D point reproducibility of 6 mm, we were able to obtain quite accurate volumes on balloons, kidneys and left ventricular molds. We feel that this accuracy was obtainable because the volumes were done on objects with fairly continuous surfaces, so the ability to exactly pinpoint the location of a point on that surface was not so critical. We also showed that, as expected, accuracy degrades as the irregularity of the objects increases. Nevertheless, all in vitro volumes were very accurate when compared with current clinical methods. We would therefore expect the data driven system to provide useful clinical estimates of organ volume, for example, left ventricular volume.

The evaluations of the system's ability to predict fetal weight have demonstrated that, as expected, accuracy degrades when going from the in vitro to the in utero situation. Nevertheless, we have shown that the addition of 3D information improves the clinical estimation of fetal weight by about 30 percent. We would expect that this accuracy could improve even more as more appropriate volume models are developed which include the limbs.

As mentioned previously the volume models developed in the current system are interpolated to whatever 3D surface data is available. They are therefore subject to noise and missing data. In addition, the system is extremely cumbersome to use because of the manual outlining of scans and the interactive nature of the volume process. The successful evaluations of the first phase of our implementation provide encouragement and medical justification for continued development of a knowledge driven system.

The research that is currently underway includes the following:

1. Completion of the survey of knowledge driven systems which could be useful for knowledge driven organ modelling.

2. Refinement of the second implementation phase: a knowledge driven modelling system which uses 3D surface points obtained with the data acquisition system developed in the first phase. This 2nd implementation phase will be designed for non-structured models and will be initially evaluated on balloons.

D. Recent Publications

[Brinkley 82a] Brinkley, J.F., McCallum, W.D., Popp, R.L. Ultrasonic three-dimensional imaging and volume: an in vitro evaluation. The American Journal of Cardiology 49:897, 1982. Abstract presented at the 31st Annual Scientific Session, American College of Cardiology, Atlanta, Georgia, April 25-29,1982.

[Brinkley 82b] Brinkley, J.F., Muramatsu, S.K., McCallum, W.D., and Popp, R.L. In vitro evaluation of an ultrasonic three-dimensional imaging and volume system. 1982. to be published in Ultrasonic Imaging.

[Brinkley 82c] Brinkley, J.F., McCallum, W.D., Muramatsu, S.K., and Liu, D.K. The Use of Ultrasonic Three-Dimensional Volume Measurements in the Estimation of Fetal Weight. In Proceedings of the 29th Annual Meeting, pages 6. Society for Gynecologic Investigation, Dallas, Texas, March, 1982.

[Brinkley 82d] Brinkley, J.F., McCallum, W.D., Muramatsu, S.K., and Liu, D.Y. Fetal weight estimation from ultrasonic three-dimensional head and trunk reconstructions: evaluation in vitro. 1982. submitted to Amer. J. Obstetrics Gynecology.

[Brinkley 82e] Brinkley, J.F., McCallum, W.D., Muramatsu, S.K., and Liu, D.Y. Fetal weight estimation from ultrasonic three-dimensional head and trunk reconstructions: evaluation in vivo. 1982. in preparation.

[Brinkley 82f] Brinkley, J.F. and McCallum, W.D. Organ volume determination by ultrasonic three-dimensional reconstruction. In Proceedings of the First AMIA Congress on Medical Informatics, pages 178-182. American Medical Informatics Association, San Francisco, May, 1982.

[McCallum 79] McCallum, W.D. and Brinkley, J.F. Estimation of fetal weight from ultrasonic measurements. Amer. J. Obstet. Gynecol. 133(2):195-200, 1979.

E. Funding Status

"Ultrasonic Three-dimensional Organ Modelling", individual postdoctoral fellowship.
Fellow: James F. Brinkley
Sponsor: W.D. McCallum
Funding Agency: National Institute of General Medical Sciences
Number: 1 F32 GM08092
Total term and direct cost: 7/1/81-6/30/84 (3 years) $55,452 (stipend)
Current funding from this fellowship: 7/1/81-6/30/82 (1 year) $17,892

II. INTERACTIONS WITH SUMEX-AIM RESOURCE

A. Collaborations

We are collaborating more with medical people than anyone else. The project is located in the Obstetrics Department at Stanford where W.D. McCallum manages the ultrasound patients. We have also been collaborating with Dr. Richard Popp in the Division of Cardiology at Stanford.

B. Sharing and Interactions with SUMEX Projects

Mostly personal contacts with the Heuristic Programming Project and MYCIN project at Stanford. The message facilities of SUMEX have been especially useful for maintaining these contacts. Since the first phase of the project is now essentially complete we have been interacting more with other SUMEX projects in order to develop the AI ideas.

C. Resource Management

In general SUMEX has been a very usable system, and the staff has been very helpful. The only complaint is that it is impossible to get anything done in the afternoons since we always get bumped.

III. RESEARCH PLANS

A. Project Goals and Plans

As mentioned in Part I we are implementing this system in phases, each phase requiring use of more sophisticated artificial intelligence techniques. The major phases are as follows (in chronological order:

1. Set up prototype system. Perform engineering tests and clinical evaluations of the ability of the system to predict fetal weight, heart and kidney volume.

As mentioned under Highlights of Research Progress, this first implementation phase is complete, as are the initial engineering and clinical evaluations. The evaluations have shown that the data driven system can provide more accurate volumes and weights than are currently clinically available. However, our patient studies have demonstrated the basic limitations of the system, which are inadequate models and difficulty of use. From a medical point of view the next phases will be attempts to remove these limitations.

2. Explore other methods for geometric modelling, AI techniques of goal directed problem solving.

In order to develop adequate models and control strategy it will be necessary to examine other AI methods of generating models and using them to guide problem solving programs. This phase of the research is now nearing completion. For this aspect of our research the SUMEX-AIM community has been especially useful.

3. Develop program, as outlined in the introduction, with several limitations--

- only a simple organ will be modelled at first, ie not the entire fetus including limbs

- the computer will still request certain scans to be retrieved from the video disk but the operator will outline them with the light pen. Since ultrasound image quality is improving so rapidly it makes sense to wait as long as possible before attempting automated border recognition. The models and control strategies developed during this phase should be useful when actual border recognition is attempted however. This version of the system should have application in left ventricular volume calculations.

4. Extend the technique to more irregular objects structured models will be developed so that the fetal limbs can be included.

5. Add image processing hardware, develop automated border recognition software.

The models developed in the last two phases will be used to guide the border recognition process. As these phases are implemented they will continue to be tested against the clinical data acquired and stored on floppy disk by the data acquisition system. In this way we can develop new ideas while continually upgrading the clinical utility of the system.

B. Justification and Requirements for Continued SUMEX Use

The goals of this project seem to be compatible with the general goals of SUMEX, i.e., to develop the uses of artificial intelligence in medicine. The problem of three-dimensional modelling is a very general one which is probably at the very heart of our ability to see. By developing a medical imaging system that models the way clinicians approach a patient we should not only develop a useful clinical tool but also explore some very fundamental problems in AI.

The availability of a large well supported facility like SUMEX has been and will continue to be very valuable as we develop and test further implementations of the system. Our current share of the SUMEX resources is adequate.

C. Needs and Plans for Other Computing Resources Beyond SUMEX-AIM

Judging from our present experience it appears that SUMEX could not handle the amount of data required for image processing on digitized ultrasound scans. This is one of the main reasons we are proposing a distributed system in which SUMEX only directs a smaller machine to do the actual number crunching. It is also one of the reasons we are postponing direct digitization until later. As microprocessors become more powerful they will be capable of acting as slaves to an intelligent SUMEX program. The AI program will direct the image processing functions of the micro so that the data is processed in an intelligent way, but SUMEX will only see the results of that processing, not the actual data. We will thus need to keep track of developments in microcomputers so that we can develop this kind of distributed system.

D. Recommendations

Since we are planning to develop a distributed system we would hope to see these kind of systems being developed by the SUMEX resource. Projects that would be of direct interest are networks (such as ETHERNET), personal computer stations, graphics displays, etc.

II.A.4 Pilot AIM Projects

The following are descriptions of the informal pilot projects currently using the AIM portion of the SUMEX-AIM resource or the Rutgers-AIM resource pending funding, and full review and authorization.

II.A.4.1 AI-COAG: Coagulation Expert Project

AI-COAG: Coagulation Expert Project

Donald A. B. Lindberg, M.D.
School of Medicine
University of Missouri-Columbia

I. SUMMARY OF RESEARCH PROGRAM

A. Project Rationale

Experiment to form a clinical consultant program based on a formal representation of the medical knowledge of the expert in human hemostasis.

B. Medical Relevance and Collaboration

Experts in hemostasis are few and are poorly distributed, often found only at University teaching hospitals and large tertiary care centers. It would be helpful to other physicians if the specialty knowledge of these experts were accessible through a computer-based consultant system.

Such a system would have medical relevance in diagnosis, management, and continuing medical education.

The team at the University of Missouri-Columbia consists of the following individuals:

Lamont W. Gaston, M.D.
Lawrence C. Kingsland III, Ph.D.
Donald A. B. Lindberg, M.D.
Anthony D. Vanker, Ph.D.
Johannes Yesus, M.D.

Dr. Gaston is a consulting hematologist and Director of the University of Missouri Health Sciences Center Coagulation Laboratory. Dr. Kingsland is a Post-doctoral Fellow whose Ph.D. is in Electrical Engineering with a concentration in Bioengineering. Dr. Lindberg is Director of the Information Science Group and Principal Investigator. Dr. Vanker is a Post-doctoral Fellow whose Ph.D. is in Physiology. Dr. Yesus is Director of the University of Missouri Health Sciences Center Blood Bank.

Subject matter expertise and the facilities of a coagulation laboratory and patient records are being provided by the University of Missouri-Columbia to build and test the consultant system. Additional testing and consultation has been obtained from Heinz Joist, M.D., Director of the Coagulation Laboratory at the St. Louis University Hospitals.

A formal research proposal was submitted to NIH on October 27, 1980. The proposal was based on initial studies performed on the SUMEX-AIM facility and at the University of Missouri-Columbia. This proposal was approved and funded for the five-year period from July, 1981 through June, 1986.

C. Highlights of Research Progress

Accomplishments:

An initial model of the coagulation consultant was created using the AGE/UNITS package on the SUMEX-AIM facility.

A feasibility test with an early textbook-level consultant model was created using EMYCIN on the SUMEX-AIM facility.

The knowledge base of these initial models was expanded and built into a new system on local Digital Equipment Corporation LSI-11 microcomputers. A laboratory interpretation system and a hemostasis history evaluation system are operational.

The laboratory interpretation system provides a differential interpretation of the results of a six-test coagulation screening battery. It has been tested on 315 cases from the files of the University of Missouri-Columbia and the St. Louis University Coagulation Laboratories. Forty-one patterns of results of the six screening tests were encountered.

The current system contains in its knowledge base interpretations for 54 combinations of the six coagulation screening tests. TMM (Tell Me More) and TMR (Tell Me Reference) knowledge sources contain additional supporting information for the user who wishes more detail.

A hemostasis history acquisition system has been added to the consultant program. Multi-level branching logic allows a totally negative history interaction in only 10-12 CRT frames, yet the system can request, tabulate and score responses to more than 200 elements of the hemostasis history.

Research in progress:

Current research is focused on the linkage of the laboratory interpretation and the hemostasis history sub-systems, with the final result the suggestion of additional coagulation laboratory tests leading to a diagnosis.

D. List of Relevant Publications

Gaston, LW; Lindberg, DAB; Vanker, AD; Kingsland, LC III: AI/COAG, a knowledge based surrogate for the human hemostasis expert. Missouri Medicine 79:6, June 1982 (in press).

Kingsland, LC III; Gaston, LW; Vanker, AD; Lindberg, DAB: A knowledge-based consultant system for problems in human hemostasis. Proceedings of the American Medical Informatics Association Congress 82 (San Francisco, May 1982). New York: Masson Publishing Co., pp. 325-9.

Lindberg, DAB; Gaston, LW; Kingsland, LC III; Vanker, AD: AI/COAG: A knowledge-based system for consultation about human hemostasis disorders; progress report. Proceedings of the Fifth Annual Symposium on Computer Applications in Medical Care (Washington DC, November 1981). New York: IEEE, pp. 253-7.

Gaston, LW; Lindberg, DAB; Kingsland, LC III; Vanker, AD: AI/COAG: A knowledge-based system for consultation about human hemostasis disorders; progress report. Proceedings of the 1981 Fall Meeting of the American Society of Clinical Pathologists and the College of American Pathologists (Las Vegas, October 1981).

Vanker, AD; Gaston, LW; Kingsland, LC III; Lindberg, DAB: Coagulation consultation system II: AI/COAG, a progress report. Presented at the Spring 1981 Meeting, National Library of Medicine Training Programs in Health Sciences and Computer Technology (Bethesda, May 1981).

Lindberg, DAB; Gaston, LW; Kingsland, LC III; Vanker, AD; Ueno, H.: A knowledge-based system for consultation about blood coagulation studies. Proceedings of the 10th Annual Conference of the Society for Computer Medicine (San Diego, September 1980).

Vanker, AD; Gaston, LW; Lindberg, DAB; Kingsland, LC III; Ueno, H.: Coagulation consultant system: interpretation of laboratory test data. Presented at the Spring 1980 Meeting, National Library of Medicine Training Programs in Health Sciences and Computer Technology (Ohio State University, May 1980).

E. Funding support

Title of Grant:
Knowledge-Based Hemostasis/Transfusion Consultant System

Principle Investigator:
Donald A. B. Lindberg, M.D.

Funding Agency:
Department of Health and Human Services,
National Heart Lung and Blood Institute.

Grant Identification Number:
DHHS 1 R01 HL 27857-01

Total Award (dates and amounts, direct costs):
July 1, 1981 to June 30, 1986 -- $433,886

Current Period (date and amount, direct costs):
July 1, 1981 to June 30, 1982 -- $85,604

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

A. Medical Collaborations and Program Dissemination via SUMEX

We have had both personal and network contacts with members of the SUMEX community during the development of the coagulation consultant system. Personal contacts at professional meetings (Society for Computer Medicine, Symposium on Computer Applications in Medical Care, International Joint Conference on Artificial Intelligence, American Medical Informatics Association, American Society of Clinical Pathologists) have been both interesting and rewarding. Network contacts allow dissemination of ideas and feedback on new work.

B. Sharing and Interactions with Other SUMEX-AIM Projects (via computing facilities, workshops, personal contacts, etc.)

In February of 1980, David Goldman, a medical student at University of Missouri-Columbia and former pre-doctoral fellow in the Information Science Group, spent a week at Stanford University. Mr. Goldman became acquainted with several of the artificial intelligence systems under development at SUMEX, and with the help of members of the SUMEX-AIM community was able to implement a simple coagulation model in EMYCIN.

In March of 1980, Dr. Haruki Ueno attended a workshop on AGE at Stanford. Dr. Ueno learned through this workshop considerable information of direct benefit to our work. He was able to return to Columbia and implement using AGE and UNITS a more extensive coagulation consultant model.

In August of 1980, Dr. Lindberg, Dr. Gaston and Dr. Vanker attended the AIM tutorial at Stanford.

Many of the AI systems of interest to us are running on the SUMEX-AIM resource. Documentation available on-line, and personal contacts facilitated by the electronic-mail capabilities at SUMEX, have been invaluable in the development of our consultant system.

The existence of the SUMEX-AIM system and the facilities it provides for the development of knowledge-based computer consultant systems, made it possible for us to achieve early in our study the degree of competence required to build our own system. SUMEX is well set up for this type of ice-breaking assistance: the AI programs are in place and running; users can study the operational systems, work with the representation structures such as those of EMYCIN and AGE, and build the understanding of their chosen research problems which enables them to strike out on their own. Feedback is always available through the electronic-message facilities. and questions or pleas for help receive prompt response.

C. Critique of Resource Management
(community facilitation, computer services, communications services, capacity, etc.)

The SUMEX-AIM professional and technical staff have been uniformly helpful and extremely competent. Busy people have given freely both of their time and their ideas. The spirit of community fostered by the SUMEX resource and its messaging facilities is extraordinary for a group of such wide geographic distribution.

The SUMEX-AIM system is well run, but sometimes shows the signs of its success in symptoms of overloading: disk space is tight, response time slows down, and files may take several minutes to open or copy. Barring occasional problems with noisy telephone lines which are not the province of the SUMEX staff, we have had few other troubles with the system.

It has been our experience that some of the SUMEX-AIM systems are sufficiently complex (and in such a state of flux and continuing development) that the designation of a particular Stanford computer scientist as a collaborator would be extremely desirable. Typically, this would involve a graduate student working on a problem suitable for dissertation research. This points up the fact that the personal component of a collaboration loses none of its importance in remote interactions: that the training, education and dissemination aspects of SUMEX should be considered as important as merely providing computational resources.

III. RESEARCH PLANS

A. Project Goals and Plans

Near-term:

We are working to link the reasoning and the results of the laboratory test interpretation subsystem with those of the hemostasis history evaluation subsystem. The outcome will be the suggestion of specific next tests to perform, leading toward diagnosis of a particular coagulation defect.

Long-range:

Additional modules will be added to the consultant system for the diagnosis of Hemophilia A and B and von Willebrand's disease, for the management of oral anticoagulant therapy, for the diagnosis of disseminated intravascular coagulation, the diagnosis of circulating anticoagulants, and the diagnosis of platelet disorders.

B. Justification and Requirements For Continued SUMEX Use

The current resources allocated to this SUMEX-AIM pilot project are sufficient for the development and dissemination of our ideas and concepts as the knowledge-based coagulation consultant model progresses. We find membership in the SUMEX community a valuable intellectual resource from which we continue to derive major benefits.

C. Needs and Plans for Other Computing Resources Beyond SUMEX-AIM

We will continue to develop on local microcomputer and minicomputer systems such portions of the coagulation consultant model as are appropriate to the smaller systems. It is both useful and important to us to know that should we need the special resources of the SUMEX-AIM facility for larger or more complex problems, they would be available to us.

D. Recommendations for Future Community and Resource Development

The concept of individual Work Stations is important for future SUMEX-AIM developments. Since we also collaborate with Rutgers using the EXPERT system, it would make more sense from our point of view if the ultimate Work Station were to be a small VAX, perhaps the VAX-11/730. This would permit support of both Lisp and Fortran systems, in contrast to the Dolphin which seems limited to Lisp.

II.B Books, Papers, and Abstracts

The publications for the various collaborative projects are summarized in their respective progress reports. Tables of these publications have been submitted separately to the Biotechnology Program Office in the special format requested. These are not reproduced here to avoid redundancy.

II.C Resource Summary Table

Detailed resource usage information is summarized starting on page 40. Tabulations of this information have been submitted separately to the Biotechnology Program Office in the special format requested. These are not reproduced here to avoid redundancy.